SDAC-TR-79-10

LEVEL II

12

AD A091910

REVISION OF THE DETECTION/ANALYSIS SYSTEM AT SDAC

W.C. Dean, A.R. Hill, P. Kovacs, J. Martin, D.M. Miller, and D.A. Whitaker
Seismic Data Analysis Center
Teledyne Geotech, 314 Montgomery Street, Alexandria Virginia 22314

26 March 1980

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

Sponsored by
The Defense Advanced Research Projects Agency (DARPA)
DARPA Order No. 2551

Monitored By
AFTAC/VSC
312 Montgomery Street, Alexandria, Virginia 22314

80 11 13 014

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

14

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: SDAC-TR-79-10

2. GOVT ACCESSION NO.: AD-A091910 9

3. RECIPIENT'S CATALOG NUMBER:

4. TITLE (and Subtitle): 6 REVISION OF THE DETECTION/ANALYSIS SYSTEM AT SDAC

5. TYPE OF REPORT & PERIOD COVERED: Technical rept.

6. PERFORMING ORG. REPORT NUMBER:

7. AUTHOR(s): 10

W. C. Dean
A. R. Hill
P. Kovacs
J. Martin
D. M. Miller
D. A. Whitaker

8. CONTRACT OR GRANT NUMBER(s): 15 F33600-79-C-0549 new DARPA Order-2551

9. PERFORMING ORGANIZATION NAME AND ADDRESS:
Teledyne Geotech
314 Montgomery Street
Alexandria, Virginia 22314

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: VT/0706

11. CONTROLLING OFFICE NAME AND ADDRESS:
Defense Advanced Research Projects Agency
Nuclear Monitoring Research Office
1400 Wilson Blvd., Arlington, Virginia 22209

12. REPORT DATE: 11 26 March 1980

13. NUMBER OF PAGES: 65

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office):
VELA Seismological Center
313 Montgomery Street
Alexandria, Virginia 22314
12 66

15. SECURITY CLASS. (of this report): Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE:

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

Author's Report Date 12/28/79

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Seismic
Computer Selection
Signal Detection
Earthquake Analysis
Earthquake Location
Software Development
Seismic Data Analysis Center
SDAC

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report presents the computational requirements for a new detection processor and regional analysis system at the Seismic Data Analysis Center (SDAC). In addition, the report considers the computational requirements to support a research environment at the SDAC which would include the detection/regional analysis capability as well as seismic research, development of new signal processing algorithms and techniques, and seismic experiments and applications of data base management techniques. Based on these requirements, various alternative computer systems are reviewed, including the existing → next page

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

408258 Du

computer facilities at the SDAC.

The existing DEC PDP 11/70 at SDAC is a good computer for supporting program development. However, we do not recommend it for on-line DP/regional input/output bandwidth, and operating system deficiencies make it considerably slower and lead to more complex programs than other available computers.

Our recommendations are conditional. If the new system is to be devoted to the DP/regional analysis service routinely, then the SEL 32/77 computer with its high input/output and computational bandwidth and its plug compatibility with all IBM peripherals at SDAC will be the most cost effective choice. Moreover, there will be no need to convert old SDAC programs since all software for this application will be new.

On the other hand, if the new computer must serve a research environment calling for on-line DP and regional analysis experiments, batch processing for seismic research, and continual software development of new processing algorithms and techniques, then the IBM 4341 is the best choice. The new IBM computer is compatible with all existing IBM peripherals at SDAC, can run virtually all existing SDAC programs without modification, and offers the most complete software development tools and versatile operating systems. The additional capital cost of this machine would be more than offset by savings in software conversions and development.

For any new computer at SDAC we recommend the addition of a floating point array processor. For time series analysis and signal processing such a unit, at a modest cost, can increase the computional bandwidth of its host computer by an order of magnitude. This excess capacity allows many experiments to be tested thoroughly which would not otherwise be done.

REVISION OF THE DETECTION/ANALYSIS SYSTEM AT SDAC

SEISMIC DATA ANALYSIS CENTER REPORT NO.: SDAC-TR-79-10

AFTAC Project Authorization No.:	VELA VT/0706/B/AFLC
Project Title:	Seismic Data Analysis Center
ARPA Order No.:	2551
Name of Contractor:	TELEDYNE GEOTECH
Contract No.:	F33600-79-C-0549
Date of Contract:	01 October 1979
Amount of Contract:	$2,845,057
Contract Expiration Date:	30 September 1980
Project Manager	Robert R. Blandford (703) 836-3882

P. O. Box 334, Alexandria, Virginia 22313

ABSTRACT

This report presents the computational requirements for a new detection processor and regional analysis system at the Seismic Data Analysis Center (SDAC). In addition, the report considers the computational requirements to support a research environment at the SDAC which would include the detection/regional analysis capability as well as seismic research, development of new signal processing algorithms and techniques, and seismic experiments and applications of data base management techniques. Based on these requirements, various alternative computer systems are reviewed, including the existing computer facilities at the SDAC.

The existing DEC PDP 11/70 at SDAC is a good computer for supporting program development. However, we do not recommend it for on-line DP/regional analysis experiments because its limited address space (16-bit machine), low input/output bandwidth, and operating system deficiencies make it considerably slower and lead to more complex programs than other available computers.

Our recommendations are conditional. If the new system is to be devoted to the DP/regional analysis service routinely, then the SEL 32/77 computer with its high input/output and computational bandwidth and its plug compatibility with all IBM peripherals at SDAC will be the most cost effective choice. Moreover, there will be no need to convert old SDAC programs since all software for this application will be new.

On the other hand, if the new computer must serve a research environment calling for on-line DP and regional analysis experiments, batch processing for seismic research, and continual software development of new processing algorithms and techniques, then the IBM 4341 is the best choice. The new IBM computer is compatible with all existing IBM peripherals at SDAC, can run virtually all existing SDAC programs without modification, and offers the most complete software development tools and versatile operating systems. The additional capital cost of this machine would be more than offset by savings in software conversions and development.

For any new computer at SDAC we recommend the addition of a floating point array processor. For time series analysis and signal processing such a unit, at a modest cost, can increase the computational bandwidth of its host computer by an order of magnitude. This excess capacity allows many experiments to be tested thoroughly which would not otherwise be done.

TABLE OF CONTENTS

ACRONYMS AND ABBREVIATIONS

A	Amplitude used with time period, T, in measuring earthquake magnitude
AA	Automatic Association
A/D	Analog-to-digital converter
ADCCP	Advanced data communications control protocol
ANSII	American national standard for information interchange
BPI	Bytes per inch to measure tape or disk recording density
BTAM-VS	Basic telecommunications access method, virtual storage; an IBM program package
CCP	Communications Control Processor
CCITT	The International Telephone & Telegraph Consultive Committee
CRT	Cathode ray tube display
DBMS	Data base management system
DEC	Digital Equipment Corporation
DMA	Direct memory access
DoD	Department of Defense
DOS	Disk operating system
DOS/VSC	Disk operating system/virtual storage
DP	Detection processor
DPS	Detection processor system, the current version running on the IBM 360/40A computer
E&S	Evans & Sutherland (corporation and/or their graphics system)
FM	Frequency modulation
FY80	Fiscal year 1980 for the U.S. government, running from October 1, 1979 to September 30, 1980
HASP	Houston asynchronous spooling program

ACRONYMS AND ABBREVIATIONS (Continued)

IBM	International Business Machines Corporation
I/O	Input/output
ips	Inches per second
ISAM	Index sequential access method, an IBM program package
K	Kilo or 10^3
KB	Kilobytes
Kb	Kilobits
L_g	L_g phase on a seismogram
M	Mega or 10^6
MPX	An operating system available on SEL computers
MULTICS	An operating system developed at MIT for Honeywell systems which features high security
NEP	Network Event Processor, the current seismic signal association/location program package running on the IBM 360/40B, DEC PDP 11/35, and E&S picture system at the SDAC
NORSAR	Norwegian seismic array
NSS	National seismic station, an automated unattended borehole seismometer which telemeters its data to a central recording site.
OS	Operating system, usually associated with the IBM version
P	The P or compressional phase on a seismogram
PS	An IBM operating system
PRIMOS	The operating system available on PRIME computers
R&D	Research and development
RSX-11M	An operating system available on the DEC PDP 11 computer family
S	The shear phase on a seismogram
S^3	Systems, Science & Software, Inc.
SAQ	The signal arrival queue in the DP/NEP system

ACRONYMS AND ABBREVIATIONS (Continued)

SDAC	Seismic Data Analysis Center, Alexandria, Va.
SEL	Systems Engineering Laboratory, Inc.
SP-Z	The short-period vertical component or data channel from a three component seismometer
STA/LTA	The ratio of the Short Term Average to Long Term Average used in the current DP
T	Time period, used with amplitude, A, in measuring earthquake magnitude
TI9900	A 16-bit microprocessor manufactured by Texas Instruments
TOTAL	A data base management system developed by Cincom Systems, Inc.
TS44	An operating system available on the IBM 360/44 computer
UNIBUS	The main data/instruction bus used on DEC computer systems.
UNIX	An operating system available on the DEC PDP 11 family of computers
USGS	The United States Geological Survey
VM/370	A virtual machine operating system developed by IBM for their 370-line of computers
VSC	The VELA Seismological Center, Alexandria, Va.

1 REVISION OF THE DETECTION/ANALYSIS SYSTEM AT SDAC

1.1 The DP Problem

This report presents the computational requirements for a new Detection Processor (DP) at the Seismic Data Analysis Center and discusses the capabilities of several commercially available computers for performing these tasks.

A new detection processor system (DP) is needed to replace the existing one at the Seismic Data Analysis Center (SDAC). The new DP must be capable of performing several functions that the current DPS is not programmed to do. These new tasks include running one or more experimental DP algorithms in parallel with the operational version. These extra DP programs might vary in detection method, detection bandwidth, and threshold or other parameters. For example, ENSCO has recently performed DP research using maximum likelihood methods, and S^3's detector uses Hilbert transform calculated on multiple bandpasses. These two processes are probably tens of times as calculation intensive as the present, highly computer-optimized STA/LTA detector. The performance of the several DP's is to be compared statistically and by the seismic analysts responsible for generating the event bulletin. Other organizations besides the SDAC will be contracted to test their DP methods. Consequently the new DP must have the capability of incorporating algorithms and subroutines from outside sources. Moreover, the DP must have the capacity to run not only on the on-line data stream but also off-line on a data base from the archives or on a data stream different (e.g. on short-period horizontal components) than those normally used in DP.

Another function proposed for the new DP system is that of post detection processing. Examples of such post processing are spectral ratio detection of spikes, processors applied to three components of data to determine back azimuths of regional and teleseismic signals, and prediction error filtering for first motion determination.

The current system handling DP is the IBM 360/40A computer which also serves as the recording computer. This computer is over 13 years old and does not have the capacity to handle a greatly expanded data stream and increased DP functions. In addition, costs for power air conditioning and maintenance for the 360/40's are high compared to those required for modern computers. Hence, following any major revision of the SDAC computer room, the IBM 360/40's are expected to be retired.

1.2 The Regional Analysis Problem

Regional analysis is the signal association. location, and parameter estimation of local, near-regional, and regional earthquakes, in addition to the teleseismic events which the Network Event Processor (NEP) has concentrated on heretofore.

Until now the Network Event Processor (NEP) has been used primarily for association of teleseismic signals and the location of teleseismic events. No concentrated effort was spent on local or near-regional events. Although such events were detected by DP and rough locations given for many of these via NEP analysis, a number of additional analysis techniques which could have yielded additional parameters were not employed. These include high and low frequency algorithms to detect locals; rotation of the horizontal components and correlation with the vertical to determine event azimuth from a station; identification of all later phases; and the various P, L_g, and S travel time curves in NEP to calculate the epicentral distance from a single site.

The NEP analysis has concentrated on the teleseismic assocation and location for several reasons. The primary one was that stations were expected to be teleseismic distances from all foreign test sites of interest. Consequently, the NEP system did not have the CRT scrolling speed or range, the data storage/retrieval capacity, or the computational speed to enable the analyst to do a thorough job on regional analysis.

In the future, local, near-regional, and regional events will be important. This change in emphasis arises because the future seismic network may have sites located such distances away from likely foreign (and USA) test sites. These sites are apt to be unattended, with satellite links for the on-line data acquisition.

To develop the hardware and software tools for such a regional analysis system, the NEP must be redesigned. The purpose of this report is to define the requirements for the new DP and the regional analysis system. These requirements must consider the expanded number of functions to be performed, the expanded input data stream, and the experimental, R&D environment the revised DP and NEP will be operating in.

An additional objective is to consider the feasibility of the DEC PDP 11/70, already available at the SDAC, serving as the new DP.

A final objective is to consider the capabilities of other new computers available commercially for handling the new DP and regional analysis tasks.

This page intentionally left blank.

2 SYSTEM DESCRIPTIONS

2.1 DP Description

The revised DP will detect seismic signals and generate a signal arrival queue (SAQ) as does the current version, but also will allow additional experimental versions of DP to run in parallel with the routine version.

The present detection processor (DP) detects seismic signals on single short-period, vertical component (SP-Z) channels or on arrays of SP-Z components at single stations. For each signal detected, the DP lists the arrival time in the SAQ, the signal turn-off time (when the amplitude falls below a detection threshold), the signal amplitude, background noise amplitude, the station channel identification, the beam number if detection is being performed on an array, and a sequence or signal index number. In the current DP the signal amplitude is measured by the STA or short term average of the SP-Z trace over 1.8 seconds. Likewise the noise amplitude is measured by the LTA or long term average over 16 STA intervals which occurred just prior to the 1.8 seconds used for the STA measurement. Detection thresholds are measured on signal-to-noise ratios or equivalently STA-to-LTA ratios.

The current DP system at SDAC receives an input bandwidth of 16.8 kilobits per second. Included in this data stream are seismic channels from six sites in Alaska; one site in Wyoming; the National Seismic Station (NSS) prototype in Tennessee; and the long period and short period channels and a detection list from NORSAR. DP runs detections routinely on seven vertical channels. At 20 samples per second and 16 bits per sample for each channel, seven SP-Z channels into a DP represent a raw data load of 2240 bits per second. An additional 25% overhead for timing and data blocking yields a total of 2800 bits per second minimum throughput for a DP.

In the future, the number of seismic sites incoming to the SDAC may increase by a factor of 3 to 6*. Thus the total input data stream from the Communications Control Processor (CCP) to the new DP could be 50 KB/sec. If the horizontal components or SP-Z arrays were to be used, the input bandwidths would be proportionally higher.

With the current STA/LTA algorithm the DP requires approximately 425 instructions per second for each channel or roughly 3000 instructions per second for the present 7 sites. This code is very efficient. To accommodate the new DP algorithms mentioned in section 1.1 the DP would require 10 to 30 times more computations per second. Thus if the number of channels

used for detection were to increase by a factor of 3, and the complexity were to increase by a factor of 10 to 30, then the new DP should possess a bandwidth of 90,000 to 270,000 instructions per second. In addition, it would be desirable for the new DP system to possess sufficient extra processing power to allow the system to process a backlog of data at several times real time. (During this time the "experimental" DP system would be inoperative.) This extra capacity will permit catch-up operations to accommodate problems in installing new and untested DP algorithms, hardware and maintenance problems with the machine, power outages, and the like. Consequently, 100,000 machine instructions per second or more does not seem to be an unreasonable requirement if the new DP is to offer the flexibility and capability desired for an expanded seismic network and the R&D environment at the SDAC. With this extra power, the DP could offer ample time per hours, or per day, for development and trouble shooting and still be able to process the routine detection load from an expanded seismic network.

The memory used by the current DP, including both program and data storage, is approximately 200 kilobytes. In future DP versions the 20 samples per second may not be decimated, the number of channels may increase by 3 times, and the complexity of the code may increase by 3 times. Also, the current version is jammed into this memory space. If memory were increased to accommodate other DP programs executing simultaneously to the on-line version, to allow spatial coherency computations, longer data windows for spectral computations, and the like, then an increase of 5 to 10 times the current DP memory space seems reasonable. Hence the new DP computer should have 1 megabytes to 2 megabytes of memory.

*Design Report for the Computer Center of an Expanded Worldwide Seismic Network; staff of Teledyne Geotech, submitted to VSC, 14 July 1978.

TABLE 2.1

Function	Instructions/function point	Instructions/(20 sps) point
decimate 20 sps to 10 sps	1	1
filter - 6 pole recursive on 10 sp's data	12 fetches 12 multiplies 12 odds 36	18
rectify on 10 sps data	1	1/2
STA short term average	1 fetch 1 add 1 add each 9 points 2 1/9	19/18
LTA long term average exponential filter recursive on each 18 pt-window	2 fetch 1 add 2 multiplies 5	5/18
Threshold logic (one each 0.6 seconds)	1 fetch 1 divide 1 compare 1 retrieve 4	4/12
Total		21.17 instructions per (20 sps) point or 423 instructions/second-channel

The table estimates the instructions needed per (20 sample per second) point for the STA/LTA detection algorithm used on each SP-Z channel in the benchmark tests. This algorithm is similar to the on-line version used at SDAC.

2.2 Regional Event Analysis

Regional event analysis, which is the association of signals and location of local through regional events, is more difficult than the network event processing (NEP) restricted to teleseisms and will require an extension of all the NEP tools.

By regional event analysis we mean an extension of the network event processing (NEP) which at the SDAC has concentrated primarily on the teleseismic. In regional event location we must associate detection signals and locate events from the local through the regional distances (5° through 20° to 30°).

The existing NEP functions include building a working data base from tape to disk; displaying the signal arrival queue (SAQ) on a CRT; associating signals by analyst or automatic association (AA) programs; displaying P-signal waveforms from various sites on a graphics CRT; filtering, aligning, and refining arrival times on these signals; locating these events; measuring signal sizes (A) and periods (T) and computing event magnitudes; identifying later phases from these events; listing the event in an earthquake bulletin giving location, time, size, stations reporting the signals, and other parameters; purging the SAQ of the detections associated with the (now) listed event; and as time and data permit, diagnostic parameter estimation (depth, M_s/m_b, etc.) to aid in discriminating explosions from earthquakes; finally, publishing and distributing (to SDAC, VSC, ARPANET mass store) the earthquake bulletin.

The tools NEP uses include a graphics CRT for displaying, scrolling, amplifying, comparing, shifting, time picking, amplitude picking, and period (dominant frequency) picking signals and seismograms. There is also an alphanumeric CRT for displaying, scrolling, editing, and purging signal arrival queues (SAQ). There are two computers, one to build data bases on disk files and provide seismological computations such as earthquake location programs, and a second one to run the refresh graphics display unit and to provide minor computations, accept analyst commands and communicate data and instructions to and from the main computer.

The current NEP system is limited in several ways which could inhibit the development of a capable regional event location system. In data display the Evans and Sutherland (E&S) graphics system is limited to 12000 short vectors per refresh cycle (1/40 second). These 12000 vectors permit displaying

8 seismic traces of 8 inches in length or 60 seconds of short-period data sampled at 10 samples per second. If 20 samples per second are used, then the traces are only 30 seconds long. Attempts to display more traces will cause the refresh period to be increased and a definite blinking of the display will result. The quality viewing region on the E&S measures 10 by 10 inches. The scrolling speed of the system is too slow due to data buffer and data transfer limits between the IBM 360/40B scientific computer and DEC PDP 11/35 graphics control computer.

There are several enhancements which the regional event location system should contain. The graphics system should be able to display traces 50% to 100% wider than those on the E&S with 90 or more seconds of data per trace. In addition, the graphics should be able to display a minimum of 8 traces without blinking. Since the state-of-the-art graphics system today can draw 50,000 or more short vectors in 1/60 second over trace widths of 17 inches (21-inch CRT's), these units could display 15 traces of 2048 points each (102.4 seconds of data at 20 sps). Scrolling speed can be increased by higher transfer rate, larger data buffers, and local memory on the graphics processor.

The current system uses approximately 450 kilobytes of memory in the IBM 360/40B computer for the waveform files and the scientific applications program (travel time tables, location programs, etc.). This memory space is overlayed into 256 kilobytes. The DEC PDP 11/35 computer feeds updating information to the E&S picture system with its graphics control processor. The PDP 11/35 uses 217 kilobytes of memory for duplicate tables, display manipulation programs, and analyst/command interface. This memory space is overlayed into 64 kilobytes of memory for executable code. The slow interface between the DEC and IBM computers, the duplication of the waveform files and tables (50 kilobytes per second), and the high levels of overlays are the main reason for the slow responses in the NEP system.

In the new regional event location system the overlayed programs, the duplication, and the computer-to-computer interchanges could be avoided by the use of a faster computer with larger memory. A graphics control processor, serving the role of the E&S picture system will be needed in any case. If this computer had 1 megabyte of memory, there would be ample room for a larger waveform file, advanced system software, a modern data base management system (DBMS), and no overlays would be necessary. Since there will be more stations

reporting, more signals and phases detected, and more working files with partially processed detections, signals, and phases, the use of more formal DBMS than has been used in the NEP heretofore may make a significant improvement in the development time of the regional event location system as well as in the case of operation of such a system by the analysts.

This page intentionally left blank.

3 SYSTEM REQUIREMENTS

3.1 Hardware Requirements

Because we anticipate experimental DP and regional analysis algorithms, the computer hardware should have speed and bandwidth capacity several times the minimum required so that ample processing power is available to accommodate hardware failures, changes in seismic network inputs, and software development time.

To accommodate the DP and regional analysis demands outlined in the previous sections, the design of the computer hardware configuration must address processing power, I/O bandwidth, memory capacity, peripheral support, and expandability.

The processing power should be sufficient to perform all DP functions faster than real time. This excess capacity will limit the effect of downtime due to maintenance, development, and system failures. DP should run a minimum of four times real time, even as other DP experiments are executing concurrently. During FY80 the detection processing will not be constrained to real time since we plan to separate it from the real time data acquisition function. In the future detection processing will involve a tiered approach involving multiple DP's with increasingly complex algorithms. Assuming that the detection process will get its input data from our current DP tapes, one has to plan for interruptions in access to the recorded data. For example, if the detector cannot run for 24 hours because of hardware failures, and if the detector runs at twice real time, then it will require 48 hours of uninterrupted access to the data in order to catch up. For these reasons, and considering that computer hardware is cheap relative to software development costs, we feel a realistic minimum is for the primary DP to process data at four times real time.

Regional event location must also perform faster than real time. The worst case situation is that requirement to analyze all data recorded in real time, which implies the analysis of seven days of data in a 40 hour period. To perform 8 hours a day for 5 days a week, analysis must be performed at 5 times real time. Ideally, a rate of 9 times real time is desired. This time frame would allow one analyst to process a week's worth of data in five four hour sessions, and would allow him to perform those housekeeping functions associated with the task. An event bulletin representing a twenty-four hour period could easily be produced every day. Table 3.1A summarizes these estimates.

TABLE 3.1A

Regional Analysis Processing Requirements

MINIMUM:

7 days * 24 hours recorded data	=	168 hours
5 working days @ 8 hours each	=	40 hours
Amount of data to be analyzed in 1 working hour	=	168/40 = 4.2 hours

Round up to 5 times real time

PREFERRED:

5 working days @ 4 hours each	=	20 hours
Amount of data to be analyzed in 1 working hour	=	168/20 = 8.4 hours

Round up to 9 times real time

Both the minimum and the preferred estimates for regional analysis processing given on Table 3.1A indicate the least amount of processing power needed for routine bulletin preparations. These estimates cover the ability to process seismic data, recorded in real time, without loss due to unavailable hardware and thus support the disaster preparedness plan.

To enhance the processing power, several hardware features may be considered. Cache memory decreases memory access time. Interleaved memory increases processor power by allowing multiple memory accesses simultaneously. A floating point processor or hardware instructions that perform single or double precision calculations dramatically decrease floating point arithmetic and number conversion. Calculations can be increased by firmware integer and trigonometric functions. Array processors are offered as add-ons from vendors or are offered as integrated units in a computer system. Several vendors also offer writable control storage modules in which the user can implement his own set of functions in firmware. A careful examination of process requirements will dictate a selection of hardware implemented capabilities.

The I/O bandwidth is directly related to the data recorded in real time and the processing rates of the seismic functions. Currently data from seven

sites is detected in real time at 16.8 kilobits per second. A conservative projection of 100 kilobits per second has been made in the Design Report for the Computer Center of an Expanded Worldwide Seismic Network. Since DP must execute in four times real time, then it must initially accept 67.2 kilobits of data per second. If multiple DP's are executing concurrently, the rate is multiplied by the number of processes. Regional analysis data rate is 151.2 kilobits or 18.9 kilobytes. A byte in this context is an 8 bit unit, not a seismic data point. See Table 3.1B for a summary of data transfer rates.

TABLE 3.1B

Transfer Rates for Seismic Data

Minimum	Current Kilobits/second		Projected Kilobits/second	
For DP	16.8 * 4	67.2	100 * 4	400
For Regional Analysis	16.8 * 5	84	100 * 5	500
		151.2Kb		900Kb
		18.9KB		112.5KB
Desired				
For DP	16.8 * 4	67.2	100 * 4	400
For Regional Analysis	16.8 * 9	151.2	100 * 9	900
		218.4Kb		1300Kb
		27.3KB		162.5KB
Worst Case				
For DP * 3	16.8 *4*3	201.6	100*4*3	1200
For Regional Analysis*2	16.8 *9*2	302.4	100*9*2	1800
		504Kb		3000Kb
		63KB		375KB

KEY: Kb refers to Kilobits
KB refers to Kilobytes

In addition to data transfer contraints, the graphics scrolling feature for regional analysis must function at 30 times real time to obtain fluid, bi-directional horizontal motion. An estimate of 100 kilobits per second is required to satisfy a request for data. Such a request requires disk access, DMA transfer to the graphics device, and an update to the refresh buffer. A more realistic transfer rate is 300 kilobits per second. This rate estimates the system overhead and data conversion that are inherent delays in the data request. Conversions that must be considered are floating point to integer

conversion, millimicron to quantum unit conversion, demultiplexing, clipping, degainranging, beaming, or rotation.

In relation to the scrolling requirement, the graphics station itself creates a transfer problem. The estimated memory space required to refresh the graphics screen is 64 kilobytes. If the graphics station has its own local memory of this amount, then the scrolling rate is estimated to be 10 data points per second or 12 kilobits per second. If the station does not have any local memory, then the main CPU must transfer the complete graphics screen for every update to the refresh buffer. The minimum acceptable refresh rate is 40 times per second. Scrolling requires an update to the refresh buffer every cycle. The transfer rate would be 2.56 megabytes per second ($64 \times 10^3 \times 40 = 2.56 \times 10^6$). If two analysis functions are executing concurrently, then the rate could be 5.12 megabytes. A consideration for expansion that would aid throughput is a multi-bus architecture. See Table 3.1C for a composite I/O bandwidth requirement.

TABLE 3.1C

Composite Best Case I/O Bandwidth

Data Rate For 3 DP's 2 *Regional Event Location*		0.3125MB
Full Refresh Every cycle for Scrolling	12Kilobytes *40 cycles/second = 0.48MB	
2 Region Stations Scrolling Concurrently		0.96MB
		1.27MB
Round Up		1.5MB
Double for Expandability		3 3MB (megabytes)

Composite Worst Case I/O Bandwidth

Data Rate For 3 DPS's 2 Regional Event Location		0.3125MB
Full Refresh Every cycle for Scrolling	64Kilobytes *40 cycles/second = 2.56MB	
2 Region Stations Scrolling Concurrently		5.12MB
		5.43MB
Round Up		6MB
Double for Expandability		12MB (megabytes)

Since the I/O bandwidth may be demanding, a large amount of real memory should be available. MOS memory is currently available on most commercial computers. In addition memory is cheap; the cost from some vendors can be as little as $15,000 for one megabyte. To minimize I/O, increase the amount of real memory available. A large amount of memory would allow data to be demultiplexed in memory. A large percentage of disk accesses which would increase the response time to satisfy a graphics request could be eliminated. The regional event/location function will require a minimum of 64 kilobyte buffers for data plus a large number of seismic tables. An estimate of 0.5 megabytes is required for a load module. Note that the current NEP system is over 0.6 megabytes in size counting all overlays. Regional event location is similar to NEP in function. To provide a responsive environment for an analyst, it is recommended that this process be memory resident. If two analysis functions are executing concurrently, one megabyte of memory is required. That amount of memory cannot be addressed by a sixteen bit register. Therefore, it is recommended that a 32 bit address capability and a minimum of two megabytes of memory are available. See Table 3.1D for a list of memory requirements.

TABLE 3.1D

Memory Requirements

Operating System and development tools	0.5MB
Regional Analysis 1	0.5MB
Regional Analysis 2	0.5MB
DP 1	0.05MB
DP 2	0.05MB
DP 3	0.05MB
	1.65MB
Expansion, Research Experiment Program Development	.35MB
	2 megabytes

An operating system addressing direct memory requires less system overhead than a system using virtual memory techniques. If the operating system has virtual memory, there are several additional considerations. If numerous processes exist, a large page size coupled with a large disk sector size of equal units would minimize paging or I/O transfer to the paging device. The amount of real memory is very important; more memory requires less paging.

If the page size is small, the number of processings large, and the amount of memory insufficient, thrashing occurs. System performance is seriously degraded.

The requirement to process faster than real time increases the data volume required at DP and regional event location initialization. The data base should be built and ready for processing before either function starts. Data tapes archived by the real time recording will always be a viable input for seismic processing. The current DP tapes will continue to provide the data for DP or regional analysis for some time. Each tape contains 5 hours of data recorded at 16.8 kilobits per second. A 5 hour data file would require 38 megabytes of information. Conceivably multiple DP's will require separate data bases. Regional analysis executing at 5 times real time will require a maximum 40 hour data base for one 8 hour session. If it executes at nine times real time, a 4 hour session requires a 36 hour maximum. A 40 hour data base requires 302.4 megabytes. Two analysts executing concurrently could request separate data bases.

Storage space must also be available for a disk base operating system. An estimate of one megabyte should be sufficient.

A theoretical total storage space requirement, including expansion for experiments or system development, is 900 megabytes. The most popular storage device is disk. To minimize contention and increase reliability, multiple drives should share the data base. If cost permits, separate disk controllers further enhance reliability. See Table 3.1E for a delineation of data volume requirements.

Several types of peripherals are required. To support this volume of data, three 300 megabyte disks should be included. For initial implementation, 1600 bpi tape should be included to allow direct input of DP tapes. In future development, the seismic research computer will become part of a local network. Data can then be transferred through the network. The real time recording device will also archive data on 6250 bpi tapes. Therefore, expandability to support this tape is desirable on the seismic research computer. If the network is broken or archived, data is requested for research projects, this feature is desirable. In any configuration, two tape drives should be considered for reliability.

Regional event location requires two alphanumeric terminals. Two analysts would require four. Four more terminals should be available for program development or experimentation. The graphic station has not been defined as yet; however, a DMA interface will be required; two stations will require two interfaces. Refer to Figure 3.1 which depicts an initial and expanded configuration.

TABLE 3.1E

Data Volume Requirements

System	
	1 MB
DP 5 Hour DP Tape	
16.8 Kilobits/seconds	
*3600 seconds/hour	
*5 hours	
302.4×10^6 bits	
1.8 bit bytes	
37.8×10^6 bytes	
	38 MB
Regional Analysis	
5*Real time/8 hour session	
= 40 hours	
= (38 MB * 8)	302.4 MB
9 * Real time/4 hour session	
= 36 hours	
16.8 kilobits/second	
*3600 seconds/hour	
*36 hours	
277×10^6 bytes	272 MB
Worst Case	
DP *3	114 MB
Regional Analysis *2 @ 5 times real time	604.8 MB
System	1.0 MB
Expansion of 25%	179.95 MB
	899.75 MB

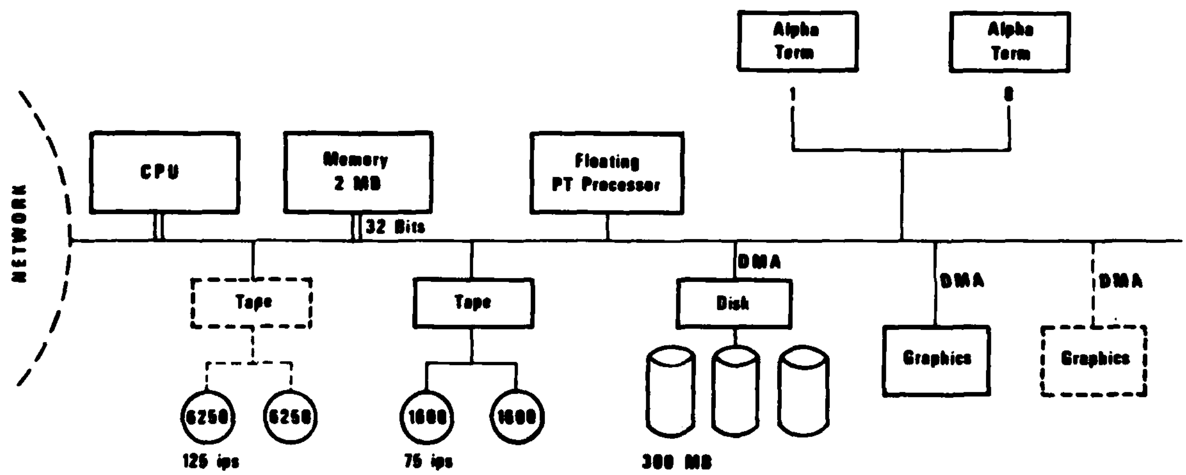

Figure 3.1 Seismic Research Environment (Dashed Figures Represent Expandable Features)

3.2 Software Requirements

System software, such as the operating system, development and debugging tools, are an important part of any new computer system; in fact, the more changeable and experimental the DP and regional analysis environment is to be, the more vital it is to have a complete set of system software packages available.

The software operating system, system services, and development tools are also an important part of the seismic research environment. Inefficient software can destroy system throughput. Software must support file access for multiple file organizations, optimizing high level languages, standard peripherals, program development tools, and real time multitasking features.

The seismic data base will consist of waveform and alphanumeric data in a mixture of files. Regional event location alone will have a minimum of 15 files open simultaneously (see Table 3.2). The definition of sequential or random accessed files has not been made. It will be determined in conjunction with throughput considerations: to save disk space or minimize I/O. Regardless of file organization, file control criteria are definable. Multiple processes or users should be able to share files. Certain files should have read-only access. Another feature is resource locking at the record level for those files that have shared access with write permission. Most vendors offer data base management facilities. In a network configuration, in which seismic functions may execute on multiple machines, a general access package is desirable so that software is portable at the user level. This feature supports the disaster preparedness plan. Depending upon the system overhead of the data base management facility, a decision to write a portable access package on top of that facility, or create a separate efficient package must be made.

The use of a high level language is imperative to support source level portability within the SDAC community or the seismic community. We recommend the use of ANSII standard FORTRAN 77 which is now offered by most vendors. The operating system should also support FORTRAN callable system services. If any modifications are to be made to the operating system itself, it too should be written in a high level language.

In support of system throughput, several operating system features are desirable: 1) asynchronous I/O at the program level; 2) ability to lock a process in memory if swapping or paging are inherent features of the operating system; 3) programmable event driven interrupts; 4) ability to share memory between processes. In addition, the ability to control the execution of

processes through command procedures allows the user to create his own environment. The NEP analyst currently uses this capability to drive the graphics task.

An important software concept is virtual memory. This feature removes the concern of memory limitations from the user. His programs can be any size desired. The operating system performs this function by dividing processes into fixed length parts called pages. A small page size will increase the amount of page checking and page faulting. A sufficient amount of real memory will minimize this system overhead.

Software reliability is as significant as hardware reliability. Even though the seismic function must execute faster than real time to avoid processing "behind" real time due to downtime, any attempt to save the system state and file modifications during power failure is desirable.

Tools for program development will aid implementation. A capable editor with character and line mode commands is imperative. A symbolic debugger is desirable. A link capability for creating executable processes from unique components aids functional modularity, and application portability.

TABLE 3.2

Regional Event Location System Files

1. Waveform File
2. Waveform Segment File
3. Detection File
4. Event File
5. Station Constant File
6. Regional Travel Time File
7. Teleseismic Travel Time File
8. Analyst Table File
9. Macro Command File
10. Error Message File
11. Scratch Pad File
12. Waveform Historical File
13. Instrument Calibration File
14. Seismic Area and Region File
15. Parameter and Initialization File

4 COMPROMISING SYSTEM REQUIREMENTS

4.1 The Costs of Compromise

Although saving capital costs by hardware compromises looks feasible at the outset of a project, the system usually costs more in both time and money to develop and fails to meet performance goals.

All of the computers being considered for the new DP and regional event location systems at SDAC are modern general purpose digital computers of roughly the same size and speed. It would seem plausible then that any of them could be programmed to do the job. Of course, given enough time and money, any of them could. However, time is crucial in two ways. First, since the DP and analysis system must keep up with a real time data acquisition system, the computer must be fast enough to complete all routine detection processing and analysis and still leave time to accommodate experiments, development time, debugging time, and maintenance problems. Second, as the programming to accommodate hardware limitations becomes more complex, the development time and debugging time for the system increases.

If the computer memory size is too small, the applications programs must be overlayed. This means that all applications programs, except for the basic root, reside on disk, and individual segments are called into memory as they are needed. In the current NEP system, the graphics task requires a total of 217 kilobytes of memory which is divided into 188 overlay segments in order to fit into the 64 kilobytes of memory available in the PDP 11/35. Such complexity makes the system hard to debug and maintain. Changes which would be minor on a simple system, such as replacement of one seismic station with another in the network, require a major revision on a complex system.

A similar situation arises when the computer uses a virtual memory operating system with a small page size. In this case, as in the case of overlayed applications programs, there is too much rolling in and out of memory of programs, data, and parameter tables. Since these pages reside on disk there is increased arm contention on the disk. With multiple users accessing the same data base and the same programs, the situation is worse. In extreme (overlay) cases or high priority (on-line analysis) situations, only one user may be permitted on the machine at a time in order to achieve any efficiency at all.

If the computer is too slow, whether due to a slow cycle time, operating system, or frequent disk access, then the system is apt to do less than we plan. If both DP and the regional analysis operations are on the same computer, then compromising on CPU speed will mean fewer DP experiments, less capability to catch up after catastrophic failures (e.g. power outages), and restrictions on the number of simultaneous users and programs. It can also lead to concentrating on the derivation of efficient algorithms and other short cuts in development rather then on having the system accomplish all the data processing tasks we would like.

By a weak operating system, we mean one with poor documentation, one difficult to understand, one with poor manufacturer support, or one with limited utilities. If the system documentation is incomplete or hard to understand, then the system development and maintenance is dependent upon system programmers more than would otherwise be necessary. This situation leads to longer development times, higher spending rates during applications developments and in-house development of system utilities that should have been provided. With poor manufacturer support, more operating system bugs must be discovered and corrected by in-house staff. Frequently, this means programming the applications around an operating system deficiency which tends to release the manufacturer from the responsibility for operating system support. Finally, all these efforts to compensate for a weak operating system makes the system development and maintenance even more dependent upon systems programmers. These specialists are highly paid, in a fluid market, and scarce. If the programmers you have at the outset of a project do not document their work completely or clearly, other systems programmers who may have to take over the project if the first ones leave have a doubly difficult time trying to find out the status of the development.

If the hardware maintenance support is poor from the manufacturer, then more reliance is placed upon in-house support. Sometimes this arrangement leads to cheaper and more responsive maintenance initially. However, as years go by, the hardware support by the manufacturer can be discontinued altogether, parts become hard to get and continuing maintenance becomes increasingly difficult. Again, this arrangement can lead to dependence upon key employees on the contractor's staff.

In summary, compromising on the system hardware and operating system requirements may seem to be feasible at the outset. It may even seem to be necessary when considering the performance goals and the budget limitations. However, these compromises frequently end up costing far more than the apparent savings in system cost, development time, and performance.

This page intentionally left blank.

5 AVAILABLE COMPUTER FACILITIES AT SDAC

5.1 The IBM 360's

Even though the IBM 360/40's and 360/44 at SDAC are old (13 years) and many peripherals, especially tape drives, are becoming harder to maintain, they are paid for, are operating a considerable investment in SDAC software, and must be expected to interface with any new computer system in both hardware and software for a while longer.

The SDAC computer room has two IBM 360/40's and an IBM 360/44 as the primary computers. Since acquiring these in 1971, SDAC has added a DEC PDP 11/70, a DEC PDP 11/35 working with an E&S graphics system, a Pluribus IMP as a Communications Control Processor (CCP) from BB&N, and a DEC PDP 15 in a nearby computer room.

The IBM 360/44 operates as the main scientific computer doing mostly batch jobs in seismic R&D with some terminals connected for program development. It has 11-1600 BPI 9-track tapes and 2-800 BPI 7-track tapes, 5-2314 (29 MB) disks, plus assorted printer, plotter, card reader, remote terminals, and ARPANET connection.

The two IBM 360/40's are used in an operational mode mostly for data acquisition, DP, and NEP/GT service. While one is on-line as a data recorder and detection processor, the other acts as back-up and serves as an NEP analysis and program development computer. In the IBM 360/40 room there are 8-1600 BPI tape units, 5-2314 (29 MB) disk drives and 4-2311 (7.2 MB) disk drives, plus a normal complement of control consoles, printer, plotter and card reader.

The PDP 11/70 has 6-1600 BPI tape drives, 1-88 MB disk and 1-300 MB disk. Most of the 16 remote terminals at SDAC are tied to the PDP 11/70 through which access can be made to the IBM 360/44. Other PDP 11/70 peripherals include a console, printer, and soon will have a new CalComp plotter.

The PDP 11/35 is an integral part of the Evans & Sutherland (E&S) graphics system on the NEP. The PDP 11/35 - E&S graphics system is coupled to the IBM 360/40B through a special interface.

The CCP is a Pluribus IMP composed of 6 Lockheed SUE computers tied to the incoming data phone lines and outputting data to the recording computer, the IBM 360/40A. The CCP has no disks or tape drives attached.

The DEC PDP-15 computer is in a separate room and not tied to any of the other IBM or DEC computers.

The IBM 360/40's and 360/44 computers are 13 years old. Most of the IBM tape drives are the same age. For most of the past 10 years these computers have operated 24 hours per day, 7 days per week. The maintenance charges for these systems amount to $110K per year. The tape drives and mechanical perpherals (typewriters, printers, card reader/punch, etc.) are requiring continually more maintenance hours to stay in service. Most of these peripherals should be replaced before very many more years pass. The IBM 360/40B has an additional 128K bytes of core added 2 years ago by Fabritek which carries lease and maintenance charges of $10.2K per year. The other IBM units, including 128 K bytes of extra core on the IBM 360/44 are purchased.

Of the other peripherals, 10-2314 type disks are purchased, 5 from Memorex and 5 from CalComp, but carry maintenance charges of $16.4K per year are paid on these units.

The DEC computers and their peripherals are purchased.

In addition to the yearly maintenance charges for IBM systems ($110K), the CalComp disks ($6.6K), Memorex disks ($9.7K), and Fabritek memory on the 40B ($10K including lease costs), there are service charges for air conditioning (7 units totalling 78 tons at $8K per year) and yearly power costs ($60K). If all IBM 360's were retired and replaced with more modern computers, the yearly charges for air conditioning service, power, and computer system maintenance could be reduced by 50% for a savings of approximately $100K per year.

There is a considerable investment in software operating on these computers. If we restrict our attention to the IBM 360's, which are the oldest computers and the one most likely to be retired, then we must consider whether any replacement computers could run this applications software without change, or if software conversions were required, how much of the software would have to be converted at what cost.

There are several large sets of applications programs running on these IBM computers. They include the DP, NEP, Library and Data Services, basic data files, and classified operational systems used to support the Air Force. We can make coarse estimates of the sizes of these software systems. The basic unit is one IBM card, essentially equivalent to one FORTRAN statement. The DP system has approximately 80,000 cards. The NEP system has approximately 160,000 cards. The library and data services has 12 large scale programs of roughly 3000 to 4000 cards each or approximately 45,000 cards; it also has 24 small programs of between 500 and 1500 cards each or approximately 25,000 cards. There are seismic support

files including frequency response and location parameters which are approximately 3000 cards plus their access routines which may number an additional 1000 cards. Finally, a large set of operational systems which run under OS on the IBM 360/44 are used to support AFTAC. These programs number perhaps 40,000 cards total of which a minimum of 20,000 would have to be converted if a new machine replaced the IBM 360/44. There are additional scientific programs plus the LASA system not counted in these estimates.

If a new machine is added to SDAC and one existing IBM 360/40 plus the 360/44 are retained, then these programs would continue to run on the old machines. If all 360's were replaced, some of the old programs probably would not be converted. A fair estimation is that 50% or more would have to be converted. The capabilities of the remainder would be discarded or picked up by new programs.

Even considering that most of this code is in FORTRAN, FORTRAN-to-FORTRAN conversions of 50% of this inventory represents a considerable effort. Although the effort will depend upon the amount of I/O connected with a program, a coarse estimate is 1 man-week per 1000 card program. Thus converting 50% of 354,000 cards might amount to 3.5 man-years of effort. At $50K per (loaded) programmer year, this cost is roughly $175K plus the loss of capability during the conversion period plus the diversion of the staff from new development efforts. These estimates could easily be high or low by a factor of two, but at least they reflect the order of magnitude of the conversion effort and the potential benefit of a new machine which can run the old programs without conversions.

TABLE 5.1

Coarse estimates of the applications software inventory at SDAC. Typically 1 card is 1 FORTRAN statement.

Program System	Number of Cards
DP	80,000
NEP	160,000
Data Services	
large programs	45,000
small programs	25,000
Data Files	4,000
OS Applications	40,000
Other Scientific Routines	?
Total	354,000 + ?

5.2 The DEC PDP 11/70

While the DEC PDP 11/70 at SDAC is an excellent machine for supporting program development from a number of terminals, there are strong indications that it cannot handle the DP and regional analysis data processing requirements and its UNIX operating system does not take advantage of the PDP 11/70 hardware capabilities or provide the software facilities for multiprocessing in a faster than real time environment.

SDAC has computer resources available for the development of new DP, regional analysis, and research experimentation. A PDP-11/70 computer with the UNIX operating system currently supports interactive data services and program development. The digitizer is to be moved from the TS44 computer to the PDP-11/70 to allow users to easily digitize large amounts of data, otherwise unavailable for automatic processing.

This computer can support some of the hardware requirements for new DP and regional event location. The PDP-11/70 supports 300 and 88 megabyte disks. There is sufficient room for multiple drives to support the data volume expected. Four dual density tape (800,1600) drives are currently available. Two additional tape drives will be available whenever the PDP-15 is dismantled. Memory is currently 0.5 megabytes of core memory. Another 1.5 megabytes of MOS memory can be added. The PDP-11/70 has a separate floating point processor, and two kilobytes of cache memory that increase the processing power. Alphanumeric terminals are available and terminal expansion capability is available. In addition to the UNIBUS, there is a separate I/O bus that is capable of 32 bit data transfers.

The UNIX operating system supports a general cross section of peripherals. Sequential and direct access file organization is available to the user. The tree structured file system removes all I/O considerations from the user. High level languages are available and the operating system itself is written in a high level language. UNIX provides a pleasant development environment with many development tools. Editors, on-line debuggers, and linking capabilities are available. The command interpreter is very effective in creating a flexible environment for users.

Version 7 of UNIX, which has not been installed at SDAC, is supposed to correct several deficiencies. ANSII standard FORTRAN 77 will be supported in the FORTRAN compiler. System services will be callable from FORTRAN and the hardware concept of separate instruction and data address space will be

supported from FORTRAN. This last feature will allow a process to contain a maximum data area of 64 kilobytes and a maximum instruction space of 64 kilobytes. A total load module can then be a maximum of 128 kilobytes. Without separate instruction and data space, the maximum size of a process is only 64 kilobytes. This feature is not available to FORTRAN on Version 6 of UNIX, which is currently installed. The installation date of Version 7 is projected to be the first of February 1980.

Unfortunately, the PDP-11/70 and UNIX contain drawbacks to the performance of new seismic research projects of the magnitude of DP and regional event location, drawbacks to supporting a flexible environment for arbitrary research projects of the same magnitude, and drawbacks to future expansion. UNIX is not vendor supported; it requires at least two system programmers dedicated to its maintenance. There is no detailed system documentation available from the vendor. There is no support for asynchronous I/O at the program level which is a requirement for regional event location. The interprocess communication obtained through pipes is not efficient and memory cannot be shared between processes. Processes cannot be locked in memory. Multiple processes can share a file; however, no file protection is available. The tree structure organization of disk files is a serious deterrent to achieving bus speed during DMA transfer for waveforms. The memory resident directory to obtain the physical location of a disk block contains two levels of indirection. Logically sequential disk blocks may be randomly located. Each disk block number must be located in the table, the arm positioned to the cylinder, a seek to the disk block, and the data accessed. If physically contiguous blocks are required to maximize throughput, a special disk driver will have to be created. This situation does not adhere to the UNIX philosophy.

Benchmarks reflect this degradation. A compute bound FORTRAN program was executed in a similar environment on the IBM 360/40, IBM 360/44, and the PDP-11/70. The 40 job was the slowest. The 11/70 job ran three times as fast as the 40 job. The 44 job ran 5.1 times as fast as the 40 job. With a mixture of I/O and compute bound functions, the IBM jobs run much faster due to I/O channels.

A test of the UNIX disk access revealed that random access and sequential access take approximately the same time. The expectation is that sequential access would be much faster. A loop of records were read from a direct access file, in which the record number was incremented by one. The program was

adjusted to generate the record number with a random number generator. This fact indicates that throughput will not be increased with sequential file organization on UNIX without special considerations.

A second test revealed throughput degradation to be directly related to the number of processes. The same FORTRAN program was triplicated to run concurrently and access the same file. The throughput decreased by a factor of three. Therefore, three DP's accessing the same data base could degrade their performance to one third.

An additional benchmark run on the PDP-11/70 was a preliminary DP that creates an SAQ file. This benchmark indicates that detection processing of 15 vertical channels would run in real time. This is for one DP process. Multiple processes would degrade this performance. We have also specified a requirement for faster than real time processing (4 times real time). If one assumes that 15 channels is adequate, then you need four times the 11/70's power to do one DP. More DP's would require another increase in performance. If the computer is also to support regional event location stations, that again dictates additional power.

An additional indication of the DP benchmark is that the 11/70's address space is a severe restriction. Secondly, it indicates that the majority of the detection processing is suspect in CPU intensive work. This dictates that to improve DP performance we need more computational power. Our recommendations to obtain the compute power are either a faster CPU and support hardware or the addition of an array processor on a reasonably fast CPU.

The statistics gathered by running whetstone benchmark programs on various computers are as follows:

TABLE 5.2

CPU	Operating System	I/O	Thousands of instr/sec Single Precision	Double
360/40	DOS	N	37.0	
11/70	UNIX	N	113.0	100.5
360/44	PS	N	188.6	93.4
HARRIS	500	N	776.4	
HARRIS	800	N	1409.4	
SEL	32/75	N	726.43	505.48
11/70	UNIX	Y	96.2	
360/44	PS	Y	166.2	

This table indicates that the PDP-11/70 suffers serious degradation when I/O is mixed with CPU intensive processing. This is attributable to the fact that it essentially has only one real I/O path to the devices that the seismic community uses for data storage, and its bandwidth is not great enough to allow it to compete with other types of computers which have multi-bus or I/O channel architectures.

To support arbitrary research projects, the resources should be available for easy implementation. The past experiences at SDAC indicate that available memory and secondary storage are important to prevent forced implementation in which code is rewritten to take less memory space or overlays are created. This method degrades performance and requires sophisticated system programmer aid. Labor and software are much more expensive than hardware. These research projects will be developed both within the SDAC/VSC environment and by other organizations. The projects will execute concurrently with routine operations and perhaps share the same data bases. Any restriction to easy implementation will require serious redirection of effort. The incorporation of the discrimination package of Systems, Science and Software in NEP is an instructive example. Memory and secondary storage to perform an experimental DP should be the minimum resources available.

Both our current DP and NEP environments are saturated. These conditions have been costly in limiting research modifications, costly in degraded system throughput, and costly in system programming support. Expansion in secondary storage is easy. A minimum projection of an additional DP, regional event location station, and communication front-end suggest a multibus architecture or multiple I/O channel support. The PDP-11/70 can have only 1 mass bus and there is a maximum of 4 loads.

The graphics stations would lock each other out if both stations and the disk controller resided on one mass bus. Two disk controllers would alleviate this problem. However, disk controllers are expensive, the graphics station would be required to use separate disks, and the bus would have a maximum load. Multi-busses are recommended. The PDP-11/70 is not expandable to two mass busses.

The following trip report from A. R. Hill gives further indication that a PDP-11/70 and the UNIX operating system are not a suitable environment for our tasks:

The purpose of the trip to USGS's Menlo Part installation was to attend meetings pertaining to the use of PDP-11/70 computers running the UNIX operating system for seismic analysis. Our intentions were to ascertain the applicability of their experiences to our own unique requirements of an 11/70 in a seismic environment.

The participants of the meeting were all in the preliminary stages of using the PDP-11/70 computers provided to them by the USGS. The goal of the USGS is to provide hardware and software that will enable selected organizations to collect seismic data from localized networks along the California coast to enable them to monitor the activity of the California faults and hopefully to share the collected information to provide a base of information that will aid future earthquake prediction efforts.

The following is a distillation of various pertinent comments and impressions that were expressed at the meetings:

- UNIX does not support real time applications.
- Disk I/O is a bottleneck.
- UNIX is better than RSX-11M.
- USGS is making large scale modifications to UNIX.
- USGS has version 7 UNIX.
- USGS has had poor support from DEC.
- USGS is subcontracting the majority of their software work.
- USGS feels they would be unable to support UNIX without the aid of a local expert who is a student at Berkeley.
- Various parties have had problems with the PDP-11/34 A/D system with FM tapes.
- USGS is rewriting seismic software for the PDP-11/70 that they have running on a MULTICS system.
- USGS is having problems with the program space limitations on the 11/70.
- USGS is installing Berkeley modifications to UNIX to aid their development.
- USGS has developed microprocessor A/D with detection capabilities (TI9900 based) that work well enough to usurp human analysts when applied to local events.
- USGS has discovered and fixed several errors in the Version 7 UNIX and in the FORTRAN 77 compiler.
- USGS is striving to use FORTRAN 77 wherever possible.
- USGS feels the 11/70 is limited in computation power and I/O power. Their implied future plans include the PDP VAX 11/780.

- USGS has under development a seismic database utility and a graphics utility. The software they demonstrated was superior in performance to Lincoln Labs software but was developed with different design goals.
- USGS felt that with a 512K 11/70 system that they would be limited to 5 simultaneous users of their software.
- USGS database software only allows one file to be open at a time. This limits performance.
- USGS is soliciting help from all sources to aid them with UNIX problems.

The overlay impression is that the SDAC application differs significantly from the USGS application; however, their software and experiences would be beneficial to our organization in support of research. I have left two tapes with the USGS and negotiated that they provide us with copies of their system at intervals that I specify.

It is also my opinion that the feelings and comments expressed at these meetings support our critique of the 11/70 and the UNIX operating system.

These are strong indications that PDP-11/70 cannot handle the processing requirements for new DP, regional event location, or experimentation. There are strong indications that the UNIX operating system does not take advantage of the PDP-11/70 hardware capabilities or provide software facilities for multiprocessing in a faster than real time environment.

6 ALTERNATIVES TO THE PDP 11/70

6.1 Alternate Computer Selections

Several commercial vendors are marketing computers that are attractive for supporting the new functions. Most modules have been newly introduced and not tested by general users. Most are price competitive with the dollar amount allocation in the FY80 contract. The following abstracts of each vendor initially considered denote features and user comments, if users were contacted. Table 6.1 offers a matrix of features for those vendors that have been contacted.

TABLE 6.1

		(UNIX) DEC PDP 11/70	DECVAX 11/780	SEL 32/77	PERKIN-ELMER 3240	HARRIS 800	PRIME 750	IBM 4341
MEMORY	Maximum Real (Bytes)	4M	8M	16M	16M	3072K	8M	16M
	Maximum User Space (Bytes)	64K	32M	512K	16M	12.5M	32M	16M
	Bandwidth (Bus)(Bytes/Sec)	6.5M	13.3M	26M	64M	7.9M	8M	
	Speed (Nano-Seconds)	600	600	600			600	
	Data Bus Width (Bits)	32	64	32	64	48	64	64
	Interleaving			Yes	No	No	Yes	
	ECC (1) (Single/Double)	C/D (2)	C/D (2)	C/D (2)	Yes (2)	C/D (2)	C/D (2)	C/D (2)
	Battery Back-Up		10 Min (Option)	70 Min (Option)	20 Min		No (UPS)	(Option)
CACHE	Amount (Bytes)	2K	8K	No			16K	8K
	Speed	200 Ns.	200 Ns.		150 Ns.	45 Ns.	80Ns.	N.A.
	Write-Through	Yes	Yes			Yes	No	
CPU	Writable Control Store (Bytes)	No	12K (No User Software)	8K	8K		No	12K
	Registers	8	16	8	128		16	16
	Index	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Indirect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Auto Increment/Decrement	Yes	Yes	Yes			Yes	Yes
	Floating Point Registers	GPR (3)	GPR (3)	GPR (3)	GPR (3)		2	GPR (3)
	Floating Point Accelerator	Yes	Yes	Yes	Yes	Yes	Yes	No
	Micro-Coded Scien. Func.	No	Yes	Yes	Yes		No	Yes
	Prefetch	Look-Ahead	Yes	Yes	Yes	Yes	Yes	Yes
I/O	Channel Programming	No	No	Yes	Yes	No	Yes (2 Mix @ 1.5M)	Yes
	DMA (Bytes)	Yes	Yes	No	Yes	Yes	Yes (32CH.)	No
	High-Speed (Bus/Channel	1 Bus @ 2M	4 Buses @ 2M	Mix 128 @ 2M/Chan.	4 Buses @ 10M	1 Bus @ 7.9M	32CH. @ 8M Total	2 Block MPX
	Low-Speed (Bus/Channel)	1 Bus @ 1.5M	1 Bus @ 1.5M	Mix 128 @ 2M/Chan.	1 Bus @ 1K	None	≤2048CH. @ 960K	1 Byte MPX
TAPE	Number Of Controllers (Max)	Buss-Slot Dependent	Mass-Bus Dependent	Many	28	31	-	16
	Drives/Controller	8	8	4	4	4	-	8
	Mix On Controller		800/1600	800/1600/6250	800/1600/6250		800/1600	800/1600/6250
	6250 BPI (Available/Speed)	FY81/75 IPS	FY81/75 IPS	Immed. /200 IPS	Immed. /		FY81/75 IPS	Immed. /200 IPS
DISK	Number Of Controllers (Max)	Buss-Slot Dependent	Bus-Slot Dependent	Many	28	4	2	16
	Drives/Controller	8	8	4	4	2	4	8
	Mix On Controller	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	User-Defined Interface	4	No	Yes, Good	Yes, Difficult		Yes, Difficult	No
	PDP 11/70 Interface	Yes	Yes	Yes	No	No	No	No
SOFTWARE	Multi-User	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Multi-Task	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Debugger	Yes	Yes (Link)	Yes	Yes	Yes	Yes	
	Shared Code And Data	No	Yes	Yes			Yes	Yes
	Virtual Memory System	No	Yes	No			Yes	Yes
	Page Size	N.A.	512 Bytes				2048 Bytes	4096 Bytes
FILE SYSTEM	File Management	Yes, Good	Yes, Good	Yes, Good			Yes, Excellent	Yes, Excellent
	Access Methods	Sequential/Random	Sequen/Random/F-11	Sequential/Random			SAM, ISAM, DAM	All
	PIOCS Availability	No	Yes (?)	Yes			Yes	Yes
	Lock At Record Level	No	Yes (RMS)	Yes (User)	Yes		Yes	Yes
	Sharable Files	Yes	Yes	Yes (User)	Yes		Yes	Yes
	DBMS Availability	Yes	Yes	Yes	Yes	Yes	Yes	Yes
REAL TIME	Ability To Lock Processes Into Memory	No	Yes	Yes	Yes			
	Interprocess Communications	Poor	Good	Good	Good	Yes	Good	Yes
NETWORK	Availability	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Protocols Supported	ARPANET	DECNET (x.25)	ADCAP/HASP	HASP	Hasp, x.25	Primenet, HASP, Telenet, x.25	IBM Features
	File System Interface (5)	No	Yes				Yes	Yes
	Special Features		Portions Of O. S. In Firmware	Programmable I/O Controllers Array Processor	IBM-Type I/O Controllers	On-Line Diagnostics, Pipelined CPU	80% of O. S. Is In Firmware	IBM Compatible Portions Of O. S. In Firmware
	Draw-Backs	Very Poor H. W. Maintenance	Very Poor H. W. Maintenance, New		New, User Support Bad	24-Bit Word	Max Of 8 Disk Drives, 6250 @ 75 IPS New	New
	Delivery		Estimated 9 Mo.	90 Days	90 Days		Immediate 650 With Field Upgrade To 750	60 Days With DoD Rating

(1) Error Checking/Correction
(2) C = Correct, D = Detect
(3) General Purpose Registers
(4) 4 Loads/Mass-Bus, Maximum Of 4 Maxi-Buses
(5) Capability To Access Remote File From A User Program As A Local File

HARRIS

PRO

- The computer bound bench mark, which we executed on the IBM and PDP computers at SDAC and the SEL 32.77, executes even faster on the Harris 500 and 800 models.
- The I/O bus on the Harris 800 CPU is 48 bits wide.

CON

- The 24 bit integer is not compatible with the application programs or data recording associated with the SDAC IBM and PDP environment. Archived data written on IBM computers must always be available for research functions and outside user requests. In addition, software previously written would have to be rewritten to run on a Harris computer. The 24 bits are also not contiguous; they are formatted as 1 parity bit and 7 data bits. Consequently, the use of a Harris computer has data and software conversion complications.

PERKINS ELMER

PRO

- 10 megabyte I/O channels
- IBM compatible data formats
- Common and read-only file access

CON

- One user indicates software support is negligible and that software functions are not as indicated in documentation.
- The PE 3240 has an initial delivery date of January, 1980. It has not been tested in the general user environment.
- Record locking is available only through the use of a separate file access package. This adds to the system overhead.

VAX

PRO

- Offers the real-time features required for regional event location.
- Offers 1 instruction look ahead.
- The UNIBUS and mass bus have microprocessor bus controllers.
- There all multiple busses over which taxing loads can be distributed.
- All operating system features are directly available to the FORTRAN user.
- Processing speed is enhanced by a large 8 kilobyte cache memory.
- Offers writeable control storage for the user, but there are no software tools offered.
- Teledyne Geotech at Garland is pleased with the VAX program development tools.

CON

- Data formats are not IBM compatible.
  This problem can be rectified in tape control hardware on the Western Peripherals products.
- The projected release of 6250 bpi tape support is the beginning of FY81. The drives are projected to be dual density (1600 and 6250 bpi) but only 75 ips (inches per second). Note that the current IBM tape drives at SDAC run at 125 ips.
- Data volume projections and additional hardware (array processor) processor, disk, or tape will tax even a VAX configuration with 4 mass busses. All tapes will interface to the mass bus. Disk and graphic stations should be distributed across two mass busses. User disks should be separate from data disks. An array processor will also interface to the mass bus. In the future, a local network interface will also interface to the mass bus with full duplex transfer rates of 200 kilobytes per second.
- ** There is a 9 month delivery delay from time of purchase.
- At a minimum, all memory pages associated with data for both DP and regional event location should be locked into memory to achieve throughput. The system overhead associated with a virtual memory operating system is still present, whether it is used or not.
- More real memory may be required to prevent thrashing in a virtual memory system than might otherwise be purchased for a system without a virtual memory facility. Since a part of the total memory will not participate in paging, the probability of thrashing in a smaller memory space increases. An easy way out, is to buy more memory.

** • DEC hardware maintenance and software support at SDAC has been atrocious over a period of years and varying machines.

- The DEC salesman has indicated that any software additions to the VMS operating system would be extremely difficult and that the source would be required. We know that a software driver for the System Industries disk must be written.

- The busses can only perform 1 transaction at a time.

- The page size is only 512 bytes as is the disk sector size. This small increment implies increases in I/O load for system use and data use. Currently, a waveform record on the IBM 40B disk is 3600 bytes so that 2 records can fit on 1 track. To write a 3600 byte record in increments of 512 bytes requires 8 transfers. There is no I/O chaining for data or I/O commands.

PRO

- User comments for hardware maintenance and system support are very favorable. Teledyne Geotech at Garland has 2 SEL 32/77's and is very pleased. They are impressed with the I/O bandwidth and the disk processor.

- Los Alamos uses 3 SEL's as high speed data switches among several large scale, state of the art computational computers. They are pleased with the I/O bandwidth and reliability.

- The SELBUS can perform polling, transaction, and acknowledgement simultaneously.

- Offers writable control store for users and all software development tools to develop and support it.

- Offer real-time features required for regional event location.

- All system features are accessible from FORTRAN.

- Special device drivers can be implemented into the operating system through job control language. The operating system source is not required.

- All memory busses have separate microprocessor controllers that attach to the SELBUS.

- The I/O device controllers are all microprocessors. The disk processor contains a 4 instruction queue for look ahead capability. This queue looks at all drives on the controller. Therefore, I/O can be performed concurrently among the drives and latency incurred between 2 I/O operations on the same disk are minimized.

- Command and data chaining aid throughput by offloading control from the CPU into the peripheral processor.

- File directories reside on the system disk, which may physically be another disk drive. This feature can be used to minimize contention.

- Disk records are stored on contiguous disk sectors. Disk files are stored by cylinder. One I/O statement can read an entire disk if desired. The disk processor manages the entire control and transfer sequence.

- SEL hardware is compatible across all models and has been released to the user community for several years.

SEL Continued

- The SELBUS and memory busses all have extremely high bandwidths.

- 6250 tape support is available now. Dual density drives (1600 and 6250 bpi) that run at 200 ips are available. The drivers are considered to be highly reliable. The 6250 recording technology is also proving, in the general user environment, to be the most reliable recording technology.

- A block multiplexor channel is available. IBM peripherals can be incorporated into a SEL configuration in this manner. Consequently, with the demise of the current SDAC IBM computers, necessary peripherals could be salvaged.

- Data formats are IBM compatible.

<u>CON</u>

- The disk sector size is only 768 bytes.

- The maximum user program module is 1 megabyte, using extended addressing. Separate instruction and data space are automatic, even for a FORTRAN user. This requirement defines a data space of 0.5 megabytes and maximum instruction space of 0.5 megabytes. Part of the instruction space, 0.128 megabytes, is dedicated to the operating system and the remaining .385 megabytes to the applications instructions. There are fullword and halfword instructions. Therefore .385 megabytes/4 = 87,000 fullword instructions or 0.385 megabytes/2 - 192,500 halfword instructions are possible. This amount of memory represents a fairly large program.

- The MPS operating system was released in October, 1979. This system is supposed to incorporate state of the art software techniques. The previous operating system RTM was designed for real time processing and is several years old.

PRIME 750

PRO

Hardware

- 32 Bit word size
- 8 Megabyte real address space
- 32 Megabyte virtual address space
- 128 32 bit hardware registers
- High speed 32 bit integer arithmetic unit
- Double precision floating point arithmetic unit
- 16K Bytes 80 nSec. cache memory
- Interleaved memory (64 bit wide transfers)
- Memory mapping and cache access interleaved
- Prefetch and decode through cache (4 instructions)
- 128 Byte segment size
- 2048 Byte page size
- Task dispatch and context switch implemented in hardware
- Very good 'IBMish' type I/O (Channel programs)
  - Direct Memory Access
    32 DMA channels Total of 8 Mbytes/sec
  - Direct Memory Control
    Up to 2048 DMC channels @ 960 KB/sec each
  - Direct Memory Transfer
    Multiple DMT's @ 2.5 Mbytes/sec each
  - Direct Memory Queues
    Hardware implemented 'ring' buffers
    Variable number/board
  - 'Burst' Mode I/O Support (64 bit wide transfers)

- Three basic line controllers
  - MDLC Sync. 4 lines @ 19.2 Kb
  - AMLC Async. 16 lines @ 9.6 Kb
  - PNC (Network) 8 Mbytes/sec
- CPU's are 'plug-in'
  - Prime 400 to Prime 500 upgrade takes ~ 20 min.
- 300 & 80 Mbyte disk sub-systems

Software

- PRIMOS (Operating System)
  - Supports up to 63 processes
  - Supports interactive users
  - Supports 'phantom' users
  - Supports batch jobs
  - Has spooling process
  - Has Procedure/Data segment sharing
  - Provides for 'locking' pages down
  - Has data communication support
    - IBM BISYNC
    - High-level Data Link Control (HDLC) (X.25)
    - Control Data 200UT
    - Univac 1004
    - ICL 7020
    - Prime DPTX (IBM 3271/3277 Display systems)

- Network access available at compiler level
- Supports 'event' processing
- Supports inter-task communications
- Has a distributed process executive available
- Compiler Support
    - FORTRAN (IV/ANSI 77)
        - Complete use of the 32 Mbyte address space
        - Generates reentrant/recursive object
        - Interface to File system, Data management
        - Options selected by user at compile time
        - Source compatible through Prime line
        - Network available
        - Supports IBM format DMA read/write
        - Extensive sort library
        - Standard FORTRAN function/subroutine library
        - Extensive string function/subroutine library
    - PL/I
    - COBOL
    - RPG II
    - Basic
    - PMA (Prime Macro Assembler)
- File System
    - Supports Direct and Sequential Access Methods
    - File can be any size that can be accommodated on the disk
    - Supports Multiple Index Data Access Method (MIDAS)
        - Implements multi-user access with record level locks
        - Available to FORTRAN
        - Provides add/delete/get/update/functions
        - Is a standard feature
    - Supports Network file access
    - CODASYL standard Database Management System available
    - A user may have up to 63 files opened at one time
    - Basic structure is 'tree', user can have lots of directories
    - Implementation of 'access; bits /read/write/update/run etc.

CON

Hardware

- Can have a maximum of 8 disk drives on a system
- Memory is expensive
    - 1/4 Megabyte is $15,000
    - 1/2 Megabyte is $25,000
    - 1 Megabyte is $40,000
- 6250 bpi Tape drive support
    - 1 year future
    - Will only be @ 75 ips

Software

- Is a 'virtual' environment
    - Even if you 'lock' pages, still have search overhead.

7 CONCLUSIONS

7.1 Recommendations

Our recommendations for revising the DP/NEP system at SDAC are conditional: if building an on-line system devoted to the DP/Regional Analysis tasks with plenty of I/O and computational power for extra inputs and extra on-line experiments is paramount, then we recommend the SEL 32/77; on the other hand, if the research environment is paramount with the system performing on-line DP/Regional Analysis tasks, plus storage/retrieval DBMS development, plus new signal processing and seismic R&D, then we recommend the IBM 4341.

Of all the computers investigated, there are four that warrant detailed comparison. These four selections are as follows:

SEL 32/77

IBM 4341

PRIME 750

DEC VAX

All other selections were judged as unacceptable for insufficient growth potential, newness in the market place, poor support, or inherent conversion problems.

These four selections are compared in Table 7.1A. The numbers reflect ratings on a scale of 1 to 4, in which 1 is the best rating and 4 is the poorest rating for that characteristic. The asterisks mark those rows of characteristics that are most important. Table 7.1B delineates a comparison of pertinent features.

Each of the top three selections (SEL, IBM, and PRIME), we judge acceptable for the DP, regional event location, or DP experimentation tasks. The VAX does not have the flexibility with which we would feel comfortable in considering future developments.

We strongly recommend the addition of a floating point array processor to any selection. Much of seismic processing (signal detection, beamforming, filtering, format conversion, spectral analysis, etc.) is well suited to an array processor. The compute power of the system will be increased by an order of magnitude. Only SEL offers an array processor as a fully integrated component of any hardware configuration. All other selections do accept an added processor. Floating Point Systems manufactures an array processor that users have added to all other selections. The cost and performance of both arrangements are comparable and therefore are not a basis for judging one selection as superior to another.

TABLE 7.1A

COMPUTER SELECTION CHARACTERISTICS

	SEL 32	PRIME 750	IBM 4341	VAX
* Cost	1	2	3	2
* Expandability				
Now	1	2	3	4
Future	4	2	1	4
* CPU Speed Now	3	2	4	1
Future	1	3	4	2
* I/O Now	1	3	1-	4
Future	1	3	1-	4
Network	3	1	2	3-
Memory	1	2	1	2-
Software	4	2	1	3
Operating System	4	2	1	3
* File				
Maintenance	3	2	1	3
Access	3	2	1	4
Program Development	4	2	3	1
Array Processor	1	2	2	2
* 6250 Tape @ 200 IPS	1	4	1	4
* System Maintenance	1	2	1	4-
Delivery	2	1	1	4

TABLE 7.1B

Specification Comparisons

Item	SEL	PRIME	IBM
CPU	Fastest	Faster	Slowest
Tapes Max.	1600BPI@200ips 6250BPI@200ips now	6250BPI@75ips in 1 year 1600BPI@75ips now	6250BPI@200ips now
Disc	300MB	300MB	300MB 570MB due 6/1/80
	@1.2MB/sec xfer rate	@1.2MB/sec xfer rate	@1.859MB/sec xfer rate
Memory Max.	16MB @$54K/MB	8MB @$40K/MB	16MB @$15K/MB
Array Processor	Built-in	Buy extra	Buy extra
Software	$10K Operating System, FORTRAN, Update Services	n.chg. FORTRAN $4.5K	$633/Mo. $38K over 5 yrs.
	$16K DBMS	$16K DBMS	$1243/mo. DBMS $74.5K over 5 yrs.
Op. Sys. Source	$16K	N.A.	N. A.

Our recommendations are conditional. They depend upon whether the government views the immediate task as most important, or as an initial step in the development of a fully integrated seismic research environment. We envision the basic building blocks of such an environment to be those depicted in Figure 7.1.

If the immediate DP and regional event location tasks are paramount, then the best choice for price/performance is the SEL 32 series. We recommend the model 32/77.

If the long range future of SDAC is most important and the scope of development is not yet clearly defined, then we recommend the IBM 4341. With IBM you obtain expandability in the product line, minimum software development across configurations, and excellent, hardware, software, and maintenance support. If IBM is too expensive, then we recommend the PRIME 750, given the constraints of expandability within an undefined environment.

The pertinent features that define each selection are amplified in the following sections.

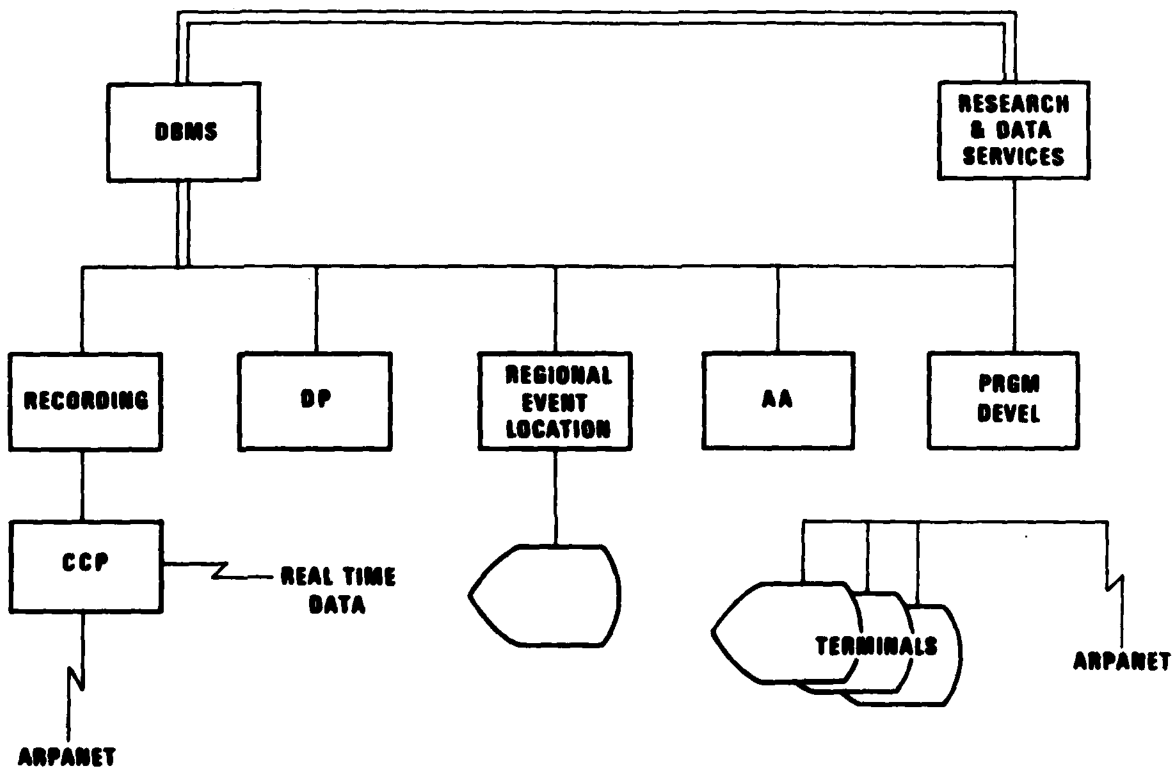

Figure 7.1 Projected Distributed SDAC R&D Environment

7.2 The SEL 32/77

If the primary role of the new computer system will be as a near real-time, operational DP and regional analysis system, then the SEL 32/77 with its high I/O and computational bandwidth, simplified RTM operating system, and low cost expandability is the first choice.

If the new computer must be primarily a DP/regional analysis machine then its role is quite specific. It will not run DP on-line, but could operate very close to real-time (within seconds) or it could play catch-up after lengthy failure, maintenance, or development shut downs. It would have extra computational power to perform all daily routine operations in a fraction of the day and to run experimental DP or regional analysis modules in parallel with the routine versions keeping statistics on the comparisons between the two. It would not be required to be a program development machine with many remote terminals throughout the SDAC. It would not serve as the DBMS machine for the entire SDAC, but would probably have a limited DBMS available for the seismic analysts in charge of generating the seismic bulletins. It would not be expected to expand into the research computer at SDAC, the role now served by the IBM 360/44.

Of those systems studied the computer which would best serve the DP/ regional analysis role is the SEL 32/77. The SEL 32 has high I/O and computational bandwidth and can easily accommodate additional inputs from an expanded seismic network. Its peripherals and peripheral controllers are plug-compatible with IBM's so that the SDAC tapes and disks could be utilized by the SEL. Its standard 1600/800 BPI tape drives are only 75 ips units. However, they do have 6250/1600 BPI, dual density drives which can operate at 200 ips. Its computation power can be increased (by 80%) at low cost ($15K) with the addition of a second computer, known as an IPU, which will automatically take on the computational tasks of the CPU but not its administrative or control tasks. Its computational power can be increased by an order of magnitude with the addition of a floating point array processor which is mounted and controlled in the SEL 32. Another SEL computer smaller than the SEL 32/77, could replace the IBM 360/40A as a recording machine. If this were done, the data acquisition, DP, and regional analysis operations would all be compatible.

Its RTM operating system is simple and is not a virtual memory system. Unpaged memory is good for a production or operational environment, espe-

cially when the input data are long single or multi-channel time series. The RTM operating system would serve the graphics function well since it is real-time oriented with standard interrupts. SEL supports HASP, the Houston Asynchronous Spooling Program, which would permit entry of jobstreams to the IBM 360's from the SEL 32.

There are several drawbacks to the choice of the SEL, most of them associated with compromising the strict DP/regional analysis role for the new computer. Because the RTM operating system is not a virtual system, the SEL 32 is not as good as others for multi-user development service. Only a few terminals could be connected reasonably to the machine. SEL has brought out MPX, a new operating system which is claimed to be able to support up to 63 users at the same time. However, being new it probably has bugs.

The maximum user memory space is one megabyte with the program instructions limited to 512 kilobytes. Its disks are sectored at 768 bytes which could also slow down access times, and hence speed, if many users were on the machine. Old SDAC programs would have to be converted to run on the SEL 32. Moreover, even though the IBM peripherals at SDAC are plug compatible with the SEL 32, IBM would probably refuse to maintain any units connected to the SEL. Either SEL or some private firm would have to maintain the IBM/SEL combination. Although acceptable service is available, maintenance of these units could become less responsive at SDAC than it has been under IBM.

Finally if the IBM 360/40B and the 360/44 computers were to be replaced, their replacement would probably not be a SEL machine. This incompatibility could lead to some difficulty in transporting seismic data and programs between the SEL and the other computers. Most manufacturers, including IBM, PRIME, and DEC are supporting the CCITT international standards (e.g. x.25, x.28, x.3) for packet switching. SEL has stated that they do not like the CCITT standards, will not support them, and instead are offering ADCCP, the Advanced Data Communications Control Protocol proposed by ANSII and used in several military networks. If the SDAC were to connect all of its machines via a local network sometime in the future, this lack of support of the CCITT standards by SEL could cause some problems.

Table 7.2 lists the major components of the SEL 32/77 system and their purchase prices and lease costs.

TABLE 7.2

SEL

1.	CPU 32/77		$46,300
	256 KB Memory	N.C.	
	768 KB Memory	40,500	
	High Speed Floating Point	7,000	
	Scientific Accel (WCS)	8,000	
	Cabinet	3,000	
	Teletype (Model 43)	2,800	
	TLC Controller	3,000	
			$64,300
			$110,600
2.	300 MB Disc & Controller		40,500
	Peripheral Cabinet		1,800
3.	Tape Drive Dual Density 75 IPS (required by diagnostics)		34,000
4.	Asynchronous Interface		3,500
			$190,400
5.	Software (RTM operating system, FORTRAN 77, other utilities)		16,165
	TOTAL (DBMS)		16,000
		One Time Charge	$32,165

Leases are available through a third party over 3 year terms. Total costs exceed 150% of the purchase price.

7.3 The IBM 4341

For a research environment calling for on-line DP and regional analysis experiments, seismic research, and continual development of new processing algorithms and techniques, the IBM 4341 is the safest choice offering the best and most complete software development tools, most versatile operating systems, and the most compatibility with current SDAC hardware and software.

If the new computer must fit into a research environment, then it must serve several different roles. It will be called upon to run on-line tests of revised DP and regional analysis systems. Moreover, these tests could include experimental DP and analysis algorithms running in parallel with the routine modules, and would also include the compilation of statistics to compare the experimental and routine versions. In addition the new computer would be called upon for continuing program and system development. Thus it should have a number of remote terminals and fairly complete software tools to support program development. The new machine should also support seismic research including the development of earthquake/explosion diagnostics. For this research the elaborate computational, data base, and graphics tools should be available to seismologists as a hands-on development system rather than being solely devoted to routine generation of daily earthquake bulletins. Finally, the system should be able to act as a development and testing system for seismic data base management tools to aid the seismic researcher as well as the analysts charged with generating seismic bulletins.

The computer serving all these roles must offer both sufficient I/O and computational bandwidth to handle the on-line or near real-time experiments and also offer many remote terminals and good program development support. The software development tools are crucial. Most of the money spent on such a system will be in software development, not initial (hardware) capital. Therefore, the more complete the software aids, and the more familiar the machine is to system and applications programmers generally, the easier (and cheaper) it will be to generate new software, and the more the SDAC staff can concentrate on applications software rather than systems software.

Of those systems studied, the computer which offers the most software support is the IBM 4341. IBM systems are the most widely known in the computer industry and so the SDAC would be less dependent upon key people remaining on the staff than with other systems. IBM systems offer the most complete software

utilities. New IBM systems will be the most compatible with the existing SDAC hardware peripherals and software. Virtually all of the programs written for the IBM 360's could be run on the IBM 4341 including those running under the old DOS operating systems. Furthermore, present IBM systems are most apt to be forward compatible with future systems, both in hardware and software. Finally, IBM offers one of the most complete data base management systems and is the firm most religious about updating all of its software packages with advanced features as the years go by.

With ample memory (2M to 4M bytes or more) and without the restrictions of small page sizes in virtual memory operating systems, the applications software would not be encumbered by multiple overlays. As a result of this larger memory and all of the software utilities, seismic applications would be simpler to program and more bug free. Thus, for the same development time and money, more seismic applications and more system parameters could be developed and tested.

The drawbacks to IBM are its higher (initial) cost. In addition, the IBM operating systems, being more versatile, represent more system overhead and so rob the system of more of its computational bandwidths than do the other computers. The apparent drawback of a long delivery time for the IBM 4300's can be overcome with a DoD priority which Teledyne already has established with IBM. As a result, delivery times for the IBM are the same as for the other systems.

Table 7.3 lists the major components of the IBM 4341 system and their purchase prices and lease costs.

TABLE 7.3

IBM

1.	CPU 4341		$245,000
	2 MB Memory	N.C.	
	Console (required)	$3,760	3,760
			$248,760
2.	300 MB Disk (3370)		35,100
	Controller (3880)		62,350
			$97,450
3.	Tape Lrive Dual Density (3420) 125 IPS		22,565
	Controller (3803)		26,300
			$48,865
4.	3274 Asynchronous Interface (max 32 lines)		18,770
			$413,845

5. Software IBM

231	VM/370 Operating System
402	Fortran
130	DOS/VSC
20	BTAM-VS
$783 *	(Free Maintenance) monthly charge

Lease costs for IBM are roughly 1/42 of the capital cost per month for a guaranteed 2 year lease, 60% applicable to the purchase.

Software is always leased with IBM for which they offer updates automatically.

7.4 The PRIME 750

If the research environment is paramount, but budget would not permit purchase of the IBM system, and if new system compatibility with old hardware and software at SDAC is not crucial, then the best compromise is the PRIME 750.

If the new computer must fit into a research environment, then it must serve the several roles of DP/regional analysis system, multi-user development system, DBMS system, and batch scientific processor which were amplified in section 7.3. However, if the constraints on the budget will not permit the purchase of the IBM system, then we must compromise among the other 32-bit machines.

Of the systems studied, the best compromise system offering the features needed for the research environment, at a price considerably under that of the IBM 4341, is the PRIME 750. Compared to the IBM, the PRIME 750 is more interactive since the design for the PRIME hardware was optimized on this feature. Its PRIMOS operating system supports up to 63 processes, interactive users, and 'phantom' users as well as batch jobs and spooling processes. PRIMOS provides for locking pages down so frequently used programs or data tables can be retained in memory. PRIME's networking software is very extensive, probably more so than IBM's. Its software costs are less than IBM's. In fact, the operating system, PRIMOS, comes free with the CPU since half of it is in firmware. Its file management system is comparable to IBM's. Actually, the lower end of the DBMS such as the Index Sequential Access Method (ISAM) and the File Management comes with the operating system just as it does with IBM's.

The PRIME 750 has the best I/O throughput of the virtual class of machines reviewed in this study, equal to if not greater than that of the IBM 4341. Because of this feature, the PRIME machines have been used widely in switching node service such as Telenet.

The maintenance on PRIME systems is good. We heard no complaints about it from users. The PRIME family includes a wide range of machines of which the 750 is the largest. There are several smaller machines which could replace the IBM 360/40A as a recording computer at SDAC. If this were done, the data acquisition, DP, and regional analysis computers would be compatible.

There are several drawbacks to the choice of the PRIME 750. Although it will be cheaper than the IBM 4341 system initially, the development costs in the long run could mean that the PRIME will be more expensive, or will have less

applications developed, or both. First of all, while the IBM machine will be compatible with the file system in NEP (ISAM, et al), neither the PRIME nor the SEL systems would be. Old disk packs cannot be read by the PRIME system directly. Moreover, converting the old NEP files on 2314 disks with ISAM, etc. is apt to carry a lower priority than other development tasks and may not be worth the time and money the conversion would require. If these files have to be converted, the IBM system could end up being cheaper than the PRIME. Since the SEL computer was chosen for work almost exclusively in the DP/regional analysis environment and not the research environment at SDAC, it may never need access to the NEP data base. In its research role, however, the PRIME machine may need the NEP file access.

The PRIME system is incompatible with the IBM peripherals at SDAC. Although the IBM tape drives are old and maintenance problems becoming more frequent, budget limitations will dictate that they not disappear suddenly or soon. Thus the PRIME system will need new tape drives, and again because of budget constraints, will probably not have very many drives attached to it initially. Another limitation is that PRIME currently offers 1600 BPI as its maximum density and those at maximum speeds of 75 ips. The 6250 BPI drives are a year away and even those are projected to appear at only 75 ips maximum speeds. Thus the tape capability of the PRIME system will be a definite step downward in capability compared to the IBM 360/40's.

The PRIME 750 is a new machine at the top of their line. Although similar in architecture to the 650 and with many subsystems in common, as of November 1979 no 750's had been delivered.

The PRIME computer can handle only 8 disk drives. However, this limitation can be extended by adding another CPU. PRIME's disk format is sectored and fixed. Only IBM offers variable format disks which provides enhanced flexibility and efficiency for SDAC seismic applications.

The memory is more expensive on the PRIME costing $40K per megabyte versus IBM's $15K.

In summary the PRIME would represent a brand new machine at SDAC, one that, unlike the SEL, is not plug compatible with SDAC peripherals nor is compatible with much of SDAC's data files and software. Such a change does entail development limitations and costs not encountered with a new IBM system.

Table 7.4 lists the major components of the PRIME 750 system and their purchase prices and lease costs.

TABLE 7.4

PRIME

1.	CPU PRIME 750		$130,000
	512 KB Memory	N.C.	
	Console (Terminet 30)	N.C.	
	1 MB Memory (additional)	$40,000	
	512 KB Memory (additional)	25,000	
2.	300 MB Disk with Controller		42,000
3.	Tape Drive 800/1600, 75 ips with Controller		19,500
4.	Asynchronous Interface (8 channel included in CPU)	N.C.	
			$191,500
5.	Software Charges		$ 4,500
	Operating System	N.C.	
	DBMS		
			$196,000

Leases are available from PRIME over a 3 year period. Total costs would exceed 150% of the purchase price.